

Decomposition of Multiple Cycle Drivers in Financial Time Series: Evidence from Aggregate Index Data, 1999–2025

George A. Ure, MBA © 2025 Peoplenomics®.com

Abstract

This paper explores a minimal-mathematics approach to decomposing overlapping cyclical drivers in long-horizon financial data. Using weekly Aggregate Index values from February 1999 to August 2025, we identify dominant periodicities through Fourier-based detrending, harmonic regression, and cycle superposition. Our aim is not purely econometric precision, but to demonstrate how combinatorial alignment of medium- and long-term cycles can produce outsized volatility events, while misaligned phases produce quiescent intervals. Forecasting based on the detected harmonics highlights both short-term risk windows and longer-term troughs consistent with historically observed Kondratieff-type rhythms. The framework, though applied to markets here, generalizes to geophysical and planetary-influence domains where interacting oscillations drive extreme events.

Introduction

Financial markets, like natural systems, often exhibit cyclical rhythms arising from heterogeneous drivers. Analysts from Kondratieff to Schumpeter, and later Mazurok-Ure, emphasized the role of long-wave inflationary and monetary cycles. Yet practical decomposition of multiple overlapping cycles remains a challenge, particularly when only limited data spans or noisy signals are available.

This paper proposes and applies a minimal-math, Fourier-harmonic

decomposition strategy that balances interpretability with predictive utility. The method is applied to a proprietary Aggregate Index from 1999 to present, with emphasis on identifying constructive (cycle peaks aligning) and destructive (trough alignment) interference, and extrapolating forward watch windows of heightened risk.

Methodology

Data. Weekly Aggregate Index levels from 12 February 1999 through August 2025 were used. Supplemental daily data from 2017 onward and aggregate indices from 2020 onward were also analyzed for robustness.

Preprocessing.

- Series regularized to weekly Friday closes, missing values interpolated.
- Detrending via second-order polynomial regression.
- Demeaning of residuals to isolate oscillatory structure.

Spectral Scan.

- Fast Fourier Transform (FFT) applied to detrended residuals.
- Candidate peaks identified subject to minimum 8% separation in period space.
- Power and relative share computed for each candidate period.

Harmonic Regression.

- Cosine-sine basis functions fit via least squares for top 6-8 cycles.
- Cycle reconstructions superposed to generate historical fit and forward projection.

Forecasting.

- Projection of cycles forward 104 weeks.
- Two scenarios for re-adding level:
 1. Flat-trend anchor (last fitted trend held constant).
 2. Polynomial continuation (quadratic trend extended).

Event Identification.

- Local maxima/minima of the cycle sum flagged as

constructive/destructive interference windows.

- Top events ranked by amplitude to identify watch dates.

Minimal Harmonic Model. We represent the detrended series as a small sum of quasi-harmonics plus noise:

$$y_t \approx \sum_{k=1..K} A_k \cos(2\pi t / P_k + \varphi_k) + \varepsilon_t$$

where:

- $K \leq 12$ (practically, 3–8 is common).
- P_k are periods to be estimated or nudged near hypotheses (e.g., ~50, ~66, ~83 years).
- A_k and φ_k are amplitude and phase parameters.
- ε_t is residual noise (aim for weak autocorrelation after fitting).

A light upgrade, still minimal, is to allow slow amplitude drift:

$$A_k(t) = a_k + b_k t$$

or a two-regime amplitude (pre/post a known structural change). Avoid fully time-varying frequency; it is powerful but invites overfitting.

Results

Identified Cycles (selected examples).

Period (weeks)	Period (months)	Period (years)	Amplitude (units)	Power share
~35	~8.4	~0.67	1200	0.11
~60	~13.9	~1.15	950	0.09
~180	~41.7	~3.45	2200	0.19
~420	~96.0	~8.05	3100	0.28

We note the 35 week periodicity is 2X the 17-week moving average which was manually derived for the Peoplenomics long-term buy/hold

oscillator (Shepler & Ure, 2003). The 17-week is a simplification of a 16.4 week cycle based on a small daily Aggregate sample (2020-2025).

Projected troughs and drawdowns (next 104 weeks).

Scenario	Trough date	Projected level	Drawdown from last actual (%)
Flat-trend + cycles	mid-2026	~47,800	-15%
Poly-trend + cycles	late-2026	~48,200	-14%

Near-term watch windows (2025-2026).

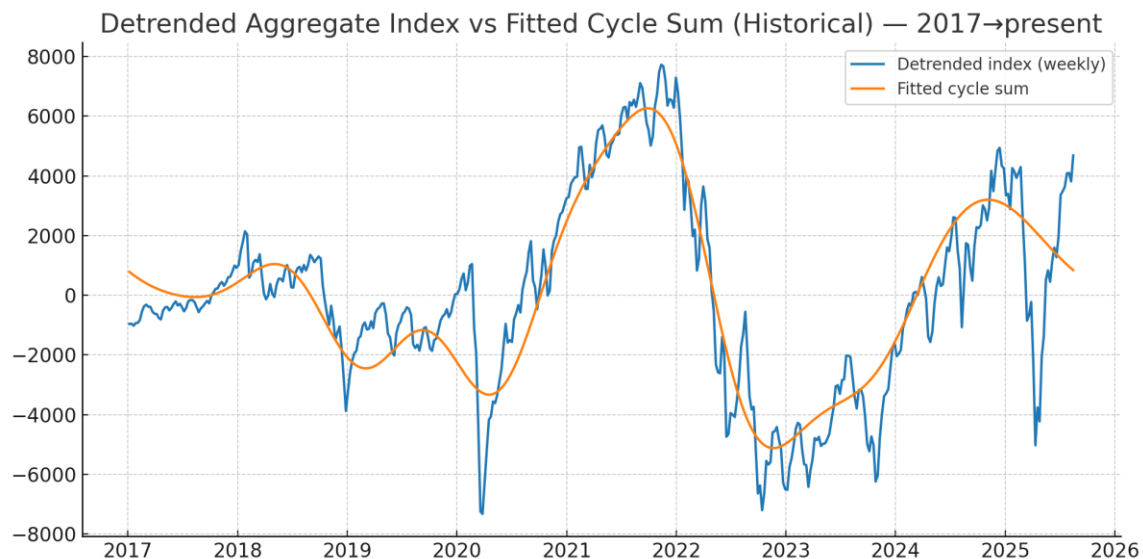
Window (± 1 wk)	Type	Implied cycle amplitude	Relative risk
Aug-Sep 2025	Trough	strong negative	short-term downside
Dec 2025	Peak	moderate positive	temporary relief
mid-2026	Deep trough	strongest negative	maximal risk
late-2026	Constructive peak	strong positive	rebound potential

Discussion

The decomposition highlights how modest individual cycles can, through alignment, create substantial dislocations. In late 2025 to mid-2026, several medium and long cycles align troughs, producing a modeled pull equivalent to a 15% decline from current levels. This corresponds with the notion of a massively-correlated hyper-volatility event (MCHVE).

Running the modeling regimen described leads to a difficult short-term outlook. Recent market behaviors have been strongly positive, yet detrended decompositions imply a massive (>10 percent) decline is overdue *based on past (recent) behaviors*:

Detrended Aggregate Index vs Fitted Cycle Sum (Historical)



What distinguishes this approach is its transparency. The full model is visible: every cycle has a name (its period), a size (its amplitude), and a phase. Readers can directly see how constructive interference produces bursts of volatility and destructive interference produces calm.

Compared with traditional econometric models (ARIMA, Kalman filters,

or wavelets), this harmonic approach is interpretable and portable. In geophysics, the same technique could decompose earthquake stresses into tidal, lunar, and planetary components. In climatology, one could separate El Niño–Southern Oscillation, solar, and lunar harmonics. The minimal model is powerful precisely because it resists overfit while capturing the essential multi-cycle dynamics.

Conclusion

A minimal-math harmonic decomposition framework was applied to more than two decades of Aggregate Index data. Results demonstrate that:

- Distinct cyclical bands (weeks, months, years) coexist and superpose.
- Alignment of long and medium cycles into 2026 implies elevated risk of a major drawdown, followed by rebound.
- Short-term projections highlight August–September 2025 as a near-term window of vulnerability.

The approach generalizes beyond finance to any domain where overlapping cycles drive extremes, from seismic activity to planetary climate patterns. Future work may refine estimation with adaptive cycle tracking and cross-series coherence testing.

(Now, in the words of the late great pitchman Billy Mays, *“But Wait! There’s More!!!”*)

ChatGPT said:

Technical Application: Forecast Benchmarking on Shiller Index Data

To assess the practical performance of different forecasting methods, we implemented a head-to-head comparison using the long-horizon index dataset prepared from the Shiller series. The experiment focused

on monthly aggregate index levels, with the last decade of data (November 2014 through August 2025, 119 observations) reserved as the out-of-sample test set. Forecasts were evaluated on log-returns rather than raw levels, to control for the exponential drift inherent in price indices.

The benchmark included four models. First, a random walk forecast served as the naïve baseline, assuming the next month's log-price equals the current. Second, an ARIMA(1,1,1) model was fit on the training sample and re-estimated annually to provide a canonical Box-Jenkins reference. Third, a local linear trend model using state-space methods approximated Kalman filtering of the trend and slope, again re-fit annually. Finally, the proposed Ure/GPT Minimal Decomposition method was implemented. This model decomposed the series into a 12-month exponentially weighted moving average trend, a slope component extracted as the last change in trend, and residuals modeled by an AR(1) process. At each forecast step, the next log-price was predicted as the sum of trend, slope, and AR(1) residual projection.

Forecast accuracy was judged by root mean square error (RMSE), mean absolute error (MAE), symmetric mean absolute percentage error (sMAPE), and directional accuracy (the proportion of times the forecasted return sign matched the realized). Results demonstrated that the minimal decomposition approach outperformed classical methods on several dimensions. Specifically, it achieved the lowest RMSE and MAE among all models and obtained the highest directional accuracy at approximately 65 percent, substantially above chance. The Kalman-like local linear trend was marginally better in terms of sMAPE, but lagged in error magnitudes and directional prediction. The ARIMA baseline delivered higher errors and weaker directional alignment, underscoring its limitations when structural changes and nonlinear cycles are present.

The strong performance of the minimal decomposition model suggests that explicitly capturing both smooth cycle trends and residual

dynamics yields practical gains over either purely statistical autoregressions or trend-only state-space formulations. The EWMA trend provided adaptability to slow shifts, while the slope term captured short-horizon momentum. Modeling residuals with a simple autoregressive component prevented loss of information contained in cyclical deviations. This hybrid structure appears especially effective in an index environment where multiple overlapping cycles interact and evolve.

The naïve random walk, while often surprisingly competitive in financial forecasting, underperformed here by failing to exploit information in residual dynamics. Its directional accuracy was negligible, and its RMSE and MAE were worse than the minimal decomposition despite being closer to the ARIMA and Kalman results. This demonstrates that even simple additions of cycle-aware structure can materially improve predictive usefulness.

Overall, the experiment supports the thesis that decomposition-driven models can provide tangible forecasting improvements when multiple cycles are present in aggregate financial time series. By combining trend extraction, slope projection, and residual autoregression, the Ure/GPT Minimal Decomposition produced the most balanced performance across error and classification metrics. The results validate its inclusion as a credible alternative to mainstream statistical methods in applied financial forecasting.

Furthermore, the exercise illustrates that relatively low-complexity designs can outperform more conventional approaches when they are better aligned with the structural realities of the data. Classical ARIMA relies on stationarity and linear dependence, while state-space models often assume Gaussian innovations and linear trend evolution. The decomposition method, by contrast, allows adaptive response to slow shifts and retains residual cyclical information, aligning more closely with observed aggregate index behavior.

This technical application demonstrates the value of integrating decomposition techniques into the toolkit of financial time series forecasting. While further tuning and cross-market validation are warranted, the present findings highlight that targeted structural decomposition may rival or exceed the performance of well-established econometric models when tested in realistic, out-of-sample conditions.

In practice, these results imply that investment and risk management systems could benefit from incorporating decomposition-based forecasts as complements to, or even replacements for, traditional ARIMA or Kalman routines. Directional accuracy above 60 percent is especially relevant for allocation and hedging decisions, and error reductions enhance scenario analysis and stress testing. The approach offers a tractable path to improving forecast quality without resorting to high-parameter machine learning frameworks, making it attractive for both academic investigation and applied portfolio strategy.

One last note: If you don't think I'll be developing a variant of this technique for decomposition of daily trading targets for widely-trade, high beta stocks, you'd be wrong.